EXHIBIT J

FOCUS - 1 of 1 DOCUMENT

Industrial Management & Data Systems

1995

SECTION: Vol. 95, No. 6 Pg. 22-26; ISSN: 0263-5577; CODEN: IMDSD8

LENGTH: 3491 words

HEADLINE: New developments in computer software

BYLINE: Thomas, Vimal; Kleiner, Brian H

BODY:

Basic concepts and principles

A brief history

The first computers were numerical calculators which seemed to contain little if any intelligence. Scientists Church and Turing held the opinion, even before the first computers were invented, that numbers were not an essential part of computation. They had an abstract conception of computation which was called symbol processing, numbers being one of the many symbols being used to represent internal states of the computing machines(1). Around the same time groups of psychologists researching the mystery of human problem solving attempted to develop programs on computers that would simulate human behaviour. Time brought together the purveyors of symbolic processing, and the psychologists concerned with human problem solving, to form a brand new interdisciplinary field of computer science that was called artificial intelligence (AI).

AI researchers have dedicated their lives to developing computer systems that produce results which we would normally associate with human intelligence. The first AI systems that were developed were crude, much like the first numerical computers that were developed.

The hardware requirements and the slowness of the AI programs made them unsuitable for any practical business use. Current advancements in electronics have produced computers with processing power which are many orders of magnitude greater than the ones in existence then. This, coupled with significant advances in the theory of AI, has brought them out of the laboratories and into the business world.

AI, like any other scientific discipline, has several distinct areas of research. One group of AI researchers is concerned primarily with developing computer programs that can read, speak or understand language as people use it in everyday conversation. These types of programs are generally referred to as natural language-processing programs. Another group of AI researchers is concerned with developing smart robots. They are especially concerned with developing the visual and tactile programs that will allow robots to observe the ongoing changes that take place as they move around in their environment(1). The third branch of AI research, the branch we are going to focus on during the cdurse of this article, is concerned with developing programs which use symbolic knowledge to simulate the behaviour of human experts. This is the branch that gave rise to knowledge-based expert systems.

Knowledge-based expert systems

To understand the concept behind knowledge-based expert systems we must first define an expert. Some define an expert as a person who knows more and more about less and less(2). Cynical as this may sound it does, however, capture the basic essence of expertise, which is specialization in a very limited domain to the exclusion of any other field that is connected to that limited domain. Thus we have an obstetrician delivering a baby, which is then handed over to a paediatrician to care for it until it is an adult, after which another group of specialists takes over depending on the type

of ailment the person suffers -- a cardiologist for the heart, a neurologist for a pinched nerve, etc. Gone are the days of a general practitioner who treated you from the cradle to the grave.

Using this definition of an expert we can define a knowledge-based expert system as an intelligent program that uses knowledge and inference procedures to solve problems that would be difficult to solve without human expertise. Like a human expert the knowledge-based expert systems are written to operate on a limited domain.

Designing and writing a system that can solve a multitude of problems over a diverse range would require a phenomenal amount of processing power and a large amount of knowledge, which would undoubtedly compromise the quality of the solutions produced by such a system. Just as a general practitioner's knowledge about the human heart cannot be relied on to diagnose and cure all forms of heart disease, knowledge of an expert system comprises facts and heuristics. The facts are a collection of information that is widely available and generally agreed on by experts in the field. Heuristics, on the other hand, are little expressed rules of good judgement, like rules of plausible reasoning and intelligent guessing that are used at expert-level decision making in the field. The people who build knowledge-based expert systems are called knowledge engineers and their technology is called knowledge engineering. AI, as a research field, is concerned primarily with studying problem solving in the abstract. Knowledge engineers, on the other hand, focus on replicating the behaviour of a specific expert when he or she is engaged in solving a narrowly defined problem. Many consider the expert system as the most practical application of artificial intelligence to date. Other areas, such as robotics, are making slower advances.

For a program to function like an expert it must be able to do things that human experts commonly do. It must ask questions to gather the data it needs, explain its conclusions, line of reasoning and why it needs the input it is asking for; and it must do all this in a language that the user of the system can easily understand. In addition, it must also be able to function with uncertain or incomplete information. A program that contains all these features may be called a true expert system.

Human information processor

It was mentioned earlier that the birth of artificial intelligence came about as a merger of the ideas of computer scientists interested in symbolic processing in computers, and psychologists involved in solving the mysteries of human problem solving. Let us look at some of the intricacies of human information processing(3,4) to gain a better understanding of the human problem-solving method, which in turn will help us understand better the computers which try to simulate human behaviour. One of the popular models of the human information processor breaks it down into three major subsystems: the perceptual subsystem, a cognitive subsystem, and a motor subsystem. The perceptual subsystem accepts inputs from the external world through its sensors, the cognitive subsystem interprets the inputs and triggers the motor subsystem to react to the input.

Sensory inputs

Stimuli from the external world are the inputs to the human information processor. These stimuli are picked up by sensors like the eyes, ears, skin and the nose. The perceptual subsystem consists of these sensors, and associated buffer memories images collected by the sensors are temporarily stored in these buffer memories before they are passed on to the cognitive subsystem for processing.

The cognitive subsystem

The cognitive subsystem consists of a working memory or a short-term memory, a long-term memory and a cognitive processor.

Working memory

Shortly after sensory information is put into the sensory buffers by the sensors of the perceptual subsystem, the cognitive processor symbolically encodes the input and puts it into the working memory of the cognitive subsystem. Not all the inputs picked up by the sensors can be put into the working memory of the cognitive subsystem since the sensors are picking up hundreds of images every second. The images are selectively encoded by the cognitive processor, a process we ordinarily refer to as "paying attention"(5).

The cognitive processor

The cognitive processor is like the central processing unit of a computer. It cycles periodically and recognizes that there is input in the sensory buffers that it needs to pick up and put in the working memory of the cognitive system. This

is called the recognize-act cycle of the cognitive processor(3) which is similar to the fetch-execute cycle of CPU of a computer. Also during each recognize-act cycle of the cognitive processor, the contents of the working memory initiate actions associatively linked to them in long-term memory.

Complex tasks involve more elaborate processing

These actions in turn modify the contents of the working memory, which in turn will initiate another action and so on until a goal is reached where no further action is required. As described above, complex tasks involve more elaborate processing and a second memory system called long-term memory

Long-term memory holds the human's mass of available knowledge(3). Some scientists believe that knowledge is stored in long-term memory in the form of symbols which are associated with one another There is a competing hypothesis which claims that the knowledge in the long-term memory is in the form of chunks which form a network of related chunks.

All the knowledge a person has about the colour green, for example, forms a chunk of knowledge. Included in this chunk are other pieces of information, such as green is a colour, leaves are green, grass is green. This chunk will be associated with another chunk of knowledge about leaves and trees and yet another chunk about grass and so on, forming a complex network which encompasses the full knowledge base of a human being.

In addition to this complex web of chunks, long-term memory is also assumed to have a complex indexing system so that, when some information is put into the working memory by the cognitive processors, information is linked associatively to a chunk of knowledge in the long-term memory and gets activated.

The contents of long-term memory are not just facts but also procedures and history about the facts. There does not appear to be any erasure of long-term memory except through destruction of brain cells; however, successful retrieval of information from the long-term memory depends on whether associations to it can be found. Retrieval can be inhibited if effective retrieval associations cannot be found or if similar associations exist to several chunks of memory causing what is commonly called "confusion". Most researchers conceptualize short-term memory as the small portion of long-term memory that is activated through association at any particular time.

There is no known limit to the amount of information that can be stored in long-term memory. The trick lies in retrieving the right kind of information from long-term memory at the right time. Sophisticated performance at a rapid rate is not uncommon among human beings, but the rapid storage of new information for long-term use is very rare.

Motor system

The motor system is the final piece of the jigsaw puzzle, that is the human information-processing system. It is translated into action by motor processors which activate voluntary muscles which in turn results in some observable activity. In the simplest tasks the cognitive system just serves to connect the perceptual system to the right outputs of the motor system. The sequence of events described above happens in performing more complex tasks that a human would normally perform.

Human problem solving

All thinking in human beings involves problem solving, regardless of how simple it may appear(G). Problem solving can be defined as finding a way to get from some initial situation to a desired goal. It usually means thinking about how to solve problems that you do not know how to solve at the onset.

There are two types of problems that we face in our day-to-day existence:

(1) well-formed problems;

(2) ill-formed problems.

In a well-formed problem we usually know where to begin, the goal that we need to reach, and the things that we need to do to reach the goal. Mathematical problems are good examples of well-formed problems. An algorithm can be used coupled with the brute force of a numerical computer to come up with a guaranteed solution. For example, if we consider all the possible moves and counter-moves that are possible in a chess game, there are about 1,000 combinations. The number of possible sequences of just eight moves would be a million billion billion. To select the best move from these combinations we have to explore billions of possibilities. No computer available today has the processing power and speed to handle such a problem.

Heuristics do not guarantee correction solutions

An ill-formed problem is of the type that may not have an explicit goal stated. It may not have a discrete set of possible solutions, and it is not very clear what you have to do to reach a solution. These are the types of problems that the human mind is usually faced with. How do you pick the winner at the Kentucky Derby? How can you pick who is going to win the Monday night football game? One needs heuristics to reach a solution for ill-formed problems. A heuristic is a method that directs thinking along the paths most likely to lead to the goal, with less promising avenues being left unexplored.

Heuristics are rules of thumb or other devices or simplifications that are used to reduce the search in problems with large search spaces. Unlike algorithms heuristics do not guarantee correct solutions. Experts who solve problems in their area of expertise usually have a lot of accumulated knowledge in that area, accumulated through years of experience in the field.

They usually acquire a knowledge of the first principles and general theories that are regarded as basic to their profession. When they practise their profession they gain experience which causes them to recompile their knowledge base (rearranging the knowledge) in their long-term memory, so that they can respond to problem situations with heuristics and theories specific to a certain area over which they then become experts.

Expert systems

We have examined how human experts solve problems by employing a large number of domain-specific facts and heuristics and the functioning of the human cognitive system. Knowledge engineers analyse the knowledge of an expert and represent that in a software that mimics the working of the human mind in the problem-solving mode.

The basic structure of an expert system contains a knowledge base similar to the long-term memory of an expert, and it contains an inference mechanism that will pick up the appropriate information from the knowledge base like the cognitive processor in the human information processor. It has a working memory which holds the results of the last thought and the input that will trigger the next thought.

In addition to these three essential components the expert system also has a knowledge acquisition subsystem that will be used to fill up the knowledge base of the system with new knowledge when it becomes available, and to correct erroneous or obsolete knowledge when necessary. It also has an explanation subsystem to provide the user of the system with an explanation of its line of reasoning, and like most computers this one has a user interface to talk to the user in a language they can understand.

Representing knowledge

The knowledge base is made up of rules and facts much like the long-term memory of an expert. There are various ways of representing the knowledge required to create a knowledge base. Representation of knowledge usually takes the form of rules operating on facts about objects of interest in the domain over which the expert system professes expertise. Delving into the details and specifics of knowledge representation would be overly technical and outside the scope of this article.

Inference engine

The second component of an expert system is the inference engine which operates on the knowledge base of the expert system in its search for a solution. The inference engine has two parts to it - the inference part and the control part. The most common inference strategy used by most inference engines is a logical rule called the modus ponens. This rule says that when A is known to be true; and if there is a rule which states, "If A then B", it is valid to conclude that B is true. That is to say that, when we know the premiss of a rule to be true, we can assume that the conclusions of the rule are justified. Modus ponens is a simple and intuitive way to conduct reasoning.

Reasoning about uncertainty

Conventional programming demands that required information be provided before computation takes place. An expert system, however, is designed to handle uncertainty much like consultants and advisers have to deal with missing or unknown information. The inference engine handles incomplete information by allowing the rules to fail when information necessary to evaluate the premisses of these rules is not available.

The second portion of the inference engine is the control portion. This portion addresses two primary problems:

(1) An expert system must have a way to decide where to start. Rules and facts reside in the knowledge base. There must be a way for the reasoning process to begin.

(2) The inference engine must resolve conflicts that occur when alternative lines of reasoning emerge. It could happen, for example, that the system reaches a point at which two or three rules can test to be true. The inference engine must choose which rule to examine next.

The issue of where to start can be addressed in two different ways depending on the nature of the expert system. If you know the question you want answered then you use a method called backward chaining. In this you start at the goal you want to achieve and work backwards through subgoals in an effort to choose an answer. If the possible outcomes are small and are known in advance, then backward chaining is very efficient.

In some cases the goal or outcome needs to be constructed or assembled, maybe because the number of possible outcomes is large. In such cases a forward- chaining strategy is used. In a forward-chaining system, premisses of the rules are examined to see whether they are true given the information on hand. If so, the conclusions are added to the facts known to be true; and the system examines the rules again.

When more than one rule tests out to be true the design of the system is usually such that it will accept the conclusions of one of the rules and ignore the conclusions of the others whose premisses are also true. The assumption being that repeated iterations will ultimately result in the right conclusions being drawn if the knowledge base is accurate.

The third component of the expert system, namely the working memory, is used to store the information collected from the user by questioning the user. This information is used to evaluate the premisses of the rules that the knowledge base contains and, depending on what premisses test out to be true, the conclusions are then returned to the working memory. The inference engine takes the conclusion that was put into the working memory and again evaluates the premisses of the rules in its knowledge base to see if a new conclusion can be drawn. When faced with the need for more information that it cannot assume from the previously drawn conclusions, it will ask the user a specific question which will allow it to proceed further, until it reaches a point where the goal is reached.

Conclusion

In conclusion, knowledge-based expert systems are -a trend of the future. When computers were first introduced to the American market they were considered to have very limited uses in areas like the Census bureau and in the scientific field. Expert systems are in a similar situation at present, but they are poised for take-off. As more and more computer vendors are recognizing the potential of this new software, more of them are making expert system shells available to the market.

Perhaps the event that signals that expert systems have come of age is the fact that IBM just introduced an expert system shell called The Integrated Reasoning Shell (TIRS). Knowing IBM's tradition of refraining from introducing a new product into the market until the technology has been perfected and the market is good and ready, one can say that knowledge-based expert systems is a technology whose time has come.

References

1. Barr, A., Cohen P.R. and Feigenbaum, E.A., The Handbook of Artificial Intelligence, Vol. 1 No. 2, William Kaufman, Los Altos, CA, 1981.

2. Nathan, G., Expert Systems and Artificial Intelligence, Weber Systems, Cleveland, OH, 1987.

3. Linsay, P.H. and Norman, D.A., Human Information Processing: An Introduction to Psychology, Academic Press, New York, NY, 1972.

4. Card, S.K., Moran, T.P. and Newell, A., The Psychology of Human-Computer Interaction, Lawrence Erlbaum Associates, Hillsdale, NJ, 1981.

5. Harmon, P. and King, D., Artificial Intelligence in Business...Expert Systems, John Wiley & Sons, New York, NY, 1985.

6. Boden, M. , Artificial Intelligence and Natural Man, Basic Books, New York, NY, 1977.

Vimal Thomas and Brian H. Kleiner are with the Department of Management, School of Business Administration and Economics, California State University, Fullerton, California, USA.

GRAPHIC: References

JOURNAL-CODE: IDS

ABI-ACC-NO: 01074205

LOAD-DATE: March 22, 1996